CONFERENCE PROCEEDINGS

RAND

Advanced Network Defense Research

Proceedings of a Workshop

Robert H. Anderson, Richard Brackney, Thomas Bozek

Prepared for the
National Security Agency

National Defense Research Institute

The conference papers described in this report were supported by the National Security Agency (NSA) and RAND's National Defense Research Institute.

ISBN: 0-8330-2938-X

The RAND conference proceedings series makes it possible to publish conference papers and discussions quickly by forgoing formal review, editing, and reformatting. Proceedings may include materials as diverse as reproductions of briefing charts, talking points, or carefully written scientific papers. Citation and quotation is permitted, but it is advisable to check with authors before citing or quoting because of the informal nature of the material.

RAND is a nonprofit institution that helps improve policy and decisionmaking through research and analysis. RAND® is a registered trademark. RAND's publications do not necessarily reflect the opinions or policies of its research sponsors.

Published 2000 by RAND
1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
RAND URL: http://www.rand.org/
To order RAND documents or to obtain additional information, contact Distribution Services: Telephone: (310) 451-7002; Fax: (310) 451-6915; Internet: order@rand.org

CONFERENCE PROCEEDINGS

RAND

Advanced Network Defense Research

Proceedings of a Workshop

Robert H. Anderson, Richard Brackney, Thomas Bozek

Prepared for the
National Security Agency

National Defense Research Institute

PREFACE

On July 11-13, 2000, approximately 20 researchers and U.S. government research sponsors met in a three-day workshop with a like number of "network defenders"--persons actively engaged in network defense analyses and measures within the Pacific Command (PACOM) and related agencies located in Oahu, Hawaii. The purpose of the workshop was a substantive exchange of information: The researchers discussed what tools, techniques, and capabilities would become available, resulting from their projects; the PACOM network defenders gave tours of their facilities (e.g., the USCINCPAC Computer Emergency Response Team (PAC-CERT) facilities, the Pacific Command Network Operations Center (NOC)), and described what alerting, analytical, and display capabilities they most needed to cope with the volume and type of information with which they deal daily.

The workshop was sponsored by the Office of the Assistant Secretary of Defense for C3I (Thomas Bozek, OASD/C3I), the US Space Command (Joseph Squatrito, USSPACECOM/J39), the Defense Advanced Research Projects Agency (Michael Skroch, DARPA/IA&S), and the National Security Agency (Richard Brackney, NSA/IADG and Larry Merritt, NSA/X).

These proceedings summarize the findings and recommendations resulting from this workshop.

For further information regarding the content of this document, please contact Richard Brackney at NSA <RCBrackney@aol.com>, Thomas Bozek at OASD/C3I (tom.bozek@osd.mil>, or Robert H. Anderson at RAND <Robert_Anderson@rand.org>.

RAND support for this workshop was provided within the Acquisition and Technology Policy Center of RAND's National Defense Research Institute (NDRI). NDRI is a federally funded research and development center sponsored by the Office of the Secretary of Defense, the Joint Staff, the defense agencies, and the unified commands.

CONTENTS

TABLES

SUMMARY

The purpose of the workshop described here was to provide an environment in which experts from industry, academe, and government could interact technically with Pacific Command (PACOM) personnel who are engaged every day in defending critical U.S. defense information systems against a variety of cyber attacks. Workshop sponsors expect to use its results to align their research objectives with the challenges facing operational users like those in the PACOM theater.

Each of the above sponsoring organizations provided plenary presentations regarding their network defense programs. Local participants then presented overviews of the operations of the PAC-CERT and the local Defense Information Systems Agency office (DISA-PAC).

Participants were provided local tours of the Regional Network Operations and Security Center (RNOSC) and the J6 Theater Command Control Cell (TCCC), to help them gain an understanding of the PACOM environment and the operating conditions of the analysts.

The workshop was structured around three breakout session themes:

1. *Analysis of Incidents* (chair: Richard Brackney, NSA)--This session identified and addressed the challenges of reducing the noise level that analysts face, obtaining tools and techniques useful in developing theater-level attack reports vs. single-point attack reports, and developing tools that will significantly reduce the time and manpower needed to analyze incident reports.
2. *Insider Misuse Mitigation* (chair: Thomas Bozek, OASD/C3I)--This session focused on assessing and obtaining user validation of the research-oriented recommendations in the recently published DoD Insider Threat Mitigation Report (DoD, 2000).
3. *Defensively Engaging the Attacker* (chair: Robert Anderson, RAND)--This session focused on how to use active network defense mechanisms, such as deception, fishbowls, lures, etc.,

in order to divert/deceive an attacker or to gain information about the attacker.

We summarize key workshop findings and recommendations under each of those three categories below.

ANALYSIS OF INCIDENTS

In intensive discussions between researchers and operators, the following R&D objectives were deemed to be of highest priority:

1. Reducing the "noise level" within the data, through automation
2. Better tools for correlating data from different sensors
3. Understanding how humans correlate intrusion detection system (IDS) information
4. Data abstraction: Creating higher-level abstractions to represent the "raw" data being received, for better human consumption and analysis
5. Means for updating the "signatures" of attacks so that they can be detected, in a manner similar to that used by virus detection tools
6. Tools for rapidly "profiling" a network: mapping its actual configuration
7. Determining what hacking tools were used in an attack
8. Maintaining the network state: "The first hack isn't enough."
9. Metrics to determine the effects of potential response scenarios.

The "Analysis of Incidents" discussion group presented a set of "R&D Findings" in five categories: those related to intrusion detection systems, network mapping/nodal analysis tools, anti-viral tools, damage assessment tools, and visualization tools.

Intrusion Detection Tools

The group felt that R&D should address and enhance current abilities to:

- *Evaluate IDSs.* Which have overlapping capabilities? Useful supplemental abilities? Which are easily tailored and configured? Which provide automated means for obtaining and installing the latest signatures? Etc.
- *Recommend which IDSs work best together.* What set of IDSs, used together, provides the best coverage? Can their outputs be easily correlated and synthesized into one overall report?
- *Disseminate information to the field.* As new capabilities are developed in IDSs, it is important that end-users "in the field" be aware of the state of the art, and that these developments are not just retained within a research community. What are the best mechanisms for providing this information dissemination on a continuing basis?
- *Develop an attack taxonomy.* Some attacks are variations on a theme of others. Some have fundamentally different approaches (e.g., a distribution denial of service, vs. a virus attack). A taxonomy of attacks should be developed, so that one can see quickly where a particular type of attack fits within the larger picture, and what the dangers and damages might be from a particular attack method.
- *Compare the attack space to IDS capabilities.* Given an attack taxonomy (see above) mapping the whole attack "space," how much of that space is covered by current IDS systems (and combinations thereof)? Where are the holes in IDS coverage?
- *Provide dynamic reallocation of IDS capabilities.* To support on-going missions, it is at times necessary to quickly reallocate IDS capabilities. Remote installation of IDSs should be possible, and the ability to install mobile software "agents" to perform various detection and analysis activities.
- *Understand how humans correlate IDS information.* In what form should IDS output data be provided, to best capitalize on human correlation capabilities? In general, we need to understand human data correlation techniques better, in order to tailor our tools to these capabilities.

- *What non-IDS information can help identify attacks?* Attacks occur in a context--of world events; of hacker tools made available on various Internet bulletin boards; of U.S. exercises, deployments, and other activities. We should understand what information external to the network and its data events can help identify the types, sources, motivations, and perpetrators of attacks.
- *Provide relevant input to researchers.* In order for researchers in IDS systems to be maximally effective, they require inputs from the field to make their studies as realistic and useful as possible.

Network Mapping/Nodal Analysis Tools

Network protection, and incident analysis, is at times hampered by a lack of understanding of the topology of the network, and an understanding of the critical nodes for its operation.

The group felt it was important that network mapping and nodal analysis tools be developed and implemented, and that network operations personnel be trained to use them.

Anti-Viral Tools

The group recommended that an "intelligent" tool be developed for opening e-mail attachments to verify that they match standard file formats. Such a tool should operate in a user-friendly safe mode.

Damage Assessment Tools

A tool is needed that will evaluate damage from viruses and from actions of unauthorized users. It should determine what the virus is doing (e.g., copying files, transmitting information), and discover what actions an unauthorized user performed (e.g., data access, modifying files, etc.)

Visualization Tools

Much better visualization tools are needed to help analysts identify patterns of activity. A rich toolkit should be available to allow output from IDSs and other sensing mechanisms to be sorted,

arranged, grouped and correlated, so that underlying patterns of interest might be detected.

INSIDER MISUSE MITIGATION

This breakout session focused on validating and prioritizing the key R&D recommendations contained in the DoD Insider Threat Mitigation Final Report of the Insider Threat Integrated Process Team (DoD-IPT 2000). Those recommendations were placed into three "bands," with the first band being of highest priority, and on which most discussion was focused. The nine DoD-IPT recommendations deemed of highest priority were the following. For each of these, the group provided a statement of:

- operational issues or constraints
- research objective
- open research issues
- success metric
- product.

The numbering on the nine "band 1" recommendations is that in the original DoD-IPT report. The one recommendation labeled "X.2" is new. Those nine highest-rated recommendations, in decreasing order of importance, were:

4.5 Create technology providing a tamper-proof (detection and resistant) audit (data including network management stuff) trail recording the actions of individuals authorized access to sensitive data and networks

6.3 Configure and deploy intrusion detection systems to monitor the activity of insiders

6.4 Implement use of network mapping tools to detect alterations in the configuration of a network

1.1 Develop and implement relevant definitions, methodologies and metrics tailored to the insider threat

4.6 Consider means by which changes can be traced in all documents generated within an organization, by simple and

tamper-proof (resistant) modifications to existing, widely used office automation programs

6.7 Develop better tools to detect, neutralize or eradicate the introduction of malicious "mobile code"

X.2 Develop a database of CND tools & their limits

1.5 Assess technologies currently available for dealing with the insider problem

4.8 (Modified) Assess the vulnerabilities of wireless connectivity and remote system/network admin both require technological approaches, with relevant R&D issues.

An additional 10 recommendations were placed in bands 2 and 3. They are listed in Section 3 of this report.

DEFENSIVELY ENGAGING THE ATTACKER

This group stated its conclusions as a set of 12 effects to be obtained in order to engage the attacker, and a set of means that should be explored in obtaining those effects. In the listing below, the effects are numbered, with the respective means shown as a bulleted list underneath. The group's effects/means findings were:

1. Deny Enemy Access

This can be achieved by:

- *"Inoculating" a system*: Provide secure, automatic transmission of patches, so that a system might be quickly "inoculated" against viruses, worms, and other attack mechanisms as they are discovered.
- *Limiting access to system and network information* (for example which operating system is being used; what network connections are in place). This can be achieved by hiding and other deceptive measures.
- *Dynamic reconfiguration of the network* to isolate an attacker.

2. Limit Enemy Access

An unauthorized party's access may be limited by:

- *Confusion (by deceptive means).* If the underlying nature and assets of an information system can be disguised to the point where an attacker is confused as to its true architecture and assets, that alone may discourage the attacker from proceeding.
- *Misinformation regarding the host.* If deliberate misinformation is provided to the attacker (e.g., regarding the type and version of the operating system in use), anyone acting on that misinformation signals a possibly hostile intent.

3. Create an Inability for an Enemy to Affect the System, Given Access

If an intruder gains some level of access to the system, we should consider active measures by which he or she is denied the ability to affect this system in deleterious ways.

4. Learn the Attacker's Intent, Identification, and Level of Sophistication

This might be achieved through the following means:

- *Entice the attacker.* If the attacker remains for a considerable period of time within the system (e.g., by investigating interesting "honeypot" files created and stored for that purpose), one might increase his or her latency within the system long enough for various tracing measures to be initiated, allowing one to learn the attacker's identity.
- *Use levels of protection or deception.* By employing levels of protection or deception, each of increasing subtlety or difficulty, one can gain a measure of the attacker's level of sophistication by discovering how many such levels he or she has broken through, or seen through.
- *Deflect or divert the attacker* to alternative sites with diverse types of information. This is a variation on the "honeypot" idea: By "pointing" the attacker at various diverse sites, one can gain an understanding of his or her intent and interests by learning which appear most appealing to him or her.

5. Detect the Presence of the Attacker

Means to detect the attacker should be quite familiar to a reader of this document. They include:

- *Automated tools to isolate anomalies.*
- *Visualization tools*, to aid a human in digesting large quantities of information, and finding subtle patterns within the data.
- *Use of "canaries" with the system*--attractive "bait" files that signal when they are accessed.
- *Watermarking of data*, and checking for its presence at various bottlenecks, gateways, or other intelligence collection measures. If data generated within a facility (or organization, or CINC, ...) were digitally watermarked in an unobservable and ineradicable manner, the presence of that watermark might be scanned for in abnormal places (e.g., on documents stored on hacker bulletin boards, in transmissions over various communication media).

6. Mapping Internal Relationships and Organizations of Intruders

Through various traffic analysis, SIGINT, and other means it may be possible to map the communications relationships of intruders, and organizations within which they operate. No explicit means were given for accomplishing this desired effect.

7. Training of Our People

It became clear during the deliberations of this session that a continuing training program is vital for network operations personnel. Such a training program should teach methodologies and principles that will have lasting value, not specific tools that may quickly become obsolescent. Means for achieving this training include:

- *Use playback tools* (e.g., based on tcpdump files) that allow repeatable data patterns and sequences to be investigated and studied.

- *Create games that are played in a "cockpit"*--not unlike video games, to engage the interest of sysadmins and operators. That cockpit might in fact be similar to, or identical to, a richer analysis environment to be provided to operators for normal network data analysis.
- *Use teams of defenders vs. teams of attackers on a training network.* Such a "red" vs. "blue" team atmosphere creates interest, and may generate novel ideas regarding attacks and defenses.
- *Consider a network of honeypots as part of a training network.* Training of operators might include the creation and distribution of attractive "honeypot" files within a network, so help gauge the interests of an attacker.

8. Reduce the Amount of Data to Be Analyzed

As mentioned in the deliberations of the "Analysis of Incidents" group, much remains to be done to provide automated aids to reduce the volume and type of network data to be analyzed in the search for anomalies.

9. Reduce the Noise in the Data

This important recommendation was discussed above in the deliberations of the "Analysis of Incidents" group.

10. Map and Understand Our Own Nets, Including Their "Backdoors"

Means for achieving this "understand thyself" dictum include:

- *Use commonly available mapping tools within DoD.* But it was remarked that relying on DoD-wide standards for data mapping causes delays and decided what those standards will be, and in their dissemination. Perhaps more local flexibility in choice of, and experimentation with, mapping tools is called for.
- *Check that our patches are installed.* Most attacks by outsiders exploit known, published vulnerabilities within operating systems and application packages. It is vital that patches to fix these flaws be made immediately after they

become available. Part of understanding our own network situation would be automated scans to determine whether all relevant patches are installed and operational.

11. Find Low-Level Probes; Separate Them from the Noise

One means for achieving this goal was highlighted in the session's presentation from Lincoln Laboratory. That is:

- *Consider the use of neural net technology.* As mentioned earlier, there are promising indications that appropriate use of neural nets might "remove the time element" and allow the connection of events (e.g., low-level probes) separated in time.

12. Monitor the Actions and Reactions of an Attacker

The primary means to achieve this effect is:

- *Create an environment in which it is possible to rapidly configure and install a comprehensive, portable "fishbowl"* within the software system. Using these facilities, all the activities of an intruder can be monitored in real time. (This, of course, requires detection of the attacker in real time, using other techniques listed in this and other sections of this report.)

OVERALL RECOMMENDATIONS

Several overall recommendations result from the workshop:

- It was agreed that researchers should work toward establishing a consistent format for reporting from intrusion detection systems. Cisco systems is interested in XML development of output from sensors. At the present, there is no mandatory output format for an IDS.
- Industry representatives stressed the need to find a means (at a significantly high level of authority) to let industry know the needs of the government.

- Several participants agreed to establish an informal exchange of information among analysts.

Overall, participants felt this was a very important exchange of information between researchers and potential users of that research. The willingness of PACOM and NCPAC operators to share information and act as a testbed for evaluating some promising research results was highlighted as very important in establishing a means of grounding network defense research in real-world requirements and exigencies.

STRUCTURE OF THIS DOCUMENT

In Section 1 we provide some background information on the workshop. Sections 2-4 are devoted to the three focus areas for the workshop (analysis of incidents, insider misuse mitigation, and defensively engaging the attacker). Section 5 contains overall conclusions and summary comments.

Appendix A contains the workshop's agenda. Appendix B lists participants and their affiliations.

ABBREVIATIONS

CINCPAC	Commander in Chief, Pacific Command
CNMT	Communication Network Modeling Tool
COTS	Commercial off-the-shelf (software or hardware systems)
DARPA	Defense Advanced Research Projects Agency
DISA	Defense Information Systems Agency
DISA-PAC	DISA Pacific Command office
GOTS	Government off-the-shelf (software or hardware systems)
IDS	Intrusion detection system
NSA	National Security Agency
OASD/C3I	Office of the Assistant Secretary of Defense for C3I (Command, Control, Communications and Intelligence)
PACOM	U.S. Pacific Command
RNOSC	Regional Network Operations and Security Center
TCCC	Theater Command Control Cell
USSPACECOM	U.S. Space Command

1. BACKGROUND

A workshop on "Advanced Network Defense" was held July 11-13, 2000 at the Ilikai Hotel, Waikiki, Oahu. It was the seventh in a series of jointly sponsored advanced network defense (AND) workshops that began in early 1997. The workshops have served as a forum to identify challenging defensive information operations (DIO) research problems and approaches to solutions. The last workshop focused on recommendations to mitigate the insider threat to information systems (cf. "DOD Insider Threat Mitigation Plan" in final review form by OASD(C3I)). That workshop resulted in the identification of nearly one hundred specific recommendations to enhance cyber-security, including numerous mid- to long-term research needs.

The purpose of the workshop reported on here is to have several key current and planned research directions in cyber-defense "validated" (confirmed as being important and useful) by the intended users of the research results. The workshop provided a setting in which a diverse group of experts from industry, academe, and government could interact directly with users--system administrators, CERT analysts and others engaged every day in protecting critical U.S. Defense information systems against a variety of attacks. Those attacks can involve insiders or outsiders. They may originate with hackers or be serious/complex attempts by determined foreign intelligence or defense organizations to obtain information, deny access to U.S. systems, or acquire protected information. This workshop provided an opportunity to select/plan research projects that are based on firsthand knowledge of the user problems/situation. It offered an excellent opportunity to identify the end user's most critical needs and for those users to provide guidance toward possible solutions. Examples of users include: System Administrators, INFOSEC Security Officers, counter-intelligence/forensics specialists, CERT Analysts.

USCINCPAC was chosen as a venue for this workshop to provide access to system administrators and others involved with network defense at USCINCPAC, as well as to remove people from their work day

interruptions. The opportunity to conduct a guided tour of USCINCPAC network defense facilities and a discussion of their activities related to active network defense also made this a very attractive location. This location allowed greater participation by analysts associated with PACOM, PACFLT PACAF, MARFORPAC, DISAPAC, and similar organizations situated nearby.

The theme of the workshop was validation by real-world users of several current/planned research directions. Researchers engaged users at USCINCPAC and collocated subordinate units who are on the front lines of cyber defense to achieve this validation, including proposed shifts in research focus that would best benefit the intended recipients of R&D results.

The workshop objectives were stated as follows:

> The workshop will have achieved its objectives if all participants come away with specific items of value: researchers will understand which of their results and proposed outcomes have greatest value to users, and what shifts in emphasis or direction are needed to achieve maximally useful outcomes; users will understand what outputs, devices, software, and processes are expected from the research community, on what approximate schedule, and how these outcomes may aid in the accomplishment of their mission. Another possible outcome might be the identification of potential users who could apply some of the research prototype tools, techniques, etc. This workshop could have follow-on workshops to review results, identify additional needs, and plan for further customer testbeds. The workshop could result in the implementation of a newsgroup where analysts of IDS's such as JIDS or ASIM could electronically provide data that drives our research program.

The workshop discussion/analysis focused on the following three areas:

1. ANALYSIS OF SECURITY INCIDENTS.

Intrusion detection systems (IDSs) and firewalls (FWs) are increasingly being installed throughout DOD and the nation's critical infrastructure. As a result, vast quantities of data are being collected--much of it representing "false positives" or otherwise irrelevant data.

This focus area asked: What is needed in the analysis of this data? Specifically:

- What are the capabilities of today's intrusion detection systems and what information is missing that would aid in analyzing the outputs from intrusion detection systems?
- What tools are needed to help analysts sift through the vast amount of reports and data produced by intrusion detection systems?
- Are processing systems that are based upon rules developed from domain knowledge experts a viable additional processing solution?
- What, if any, correlation is done between intrusion sensors and how? Is it mostly cognitive (i.e., human, rather than computer-algorithmic). If it is cognitive, what processes could potentially be automated and how?
- In addition to IDS database, what other sources of information can be used to assess the cyber situation?
- How would analysts like to get their information in the future--Amount, format timeliness... ?
- What is a typical successful on-line and forensic analysis and why?
- What tools are needed to minimize the resources needed to detect routine attacks? What tools are needed to help analysts focus on new/complex attacks that may occur over time and space?
- What information is important to SysAdmin's, CERT Analysts and other cyber specialists in their analysis? What information is missing that would aid in their analysis? How do they get the cyber situation information now?
- What is being done to differentiate an inside attack by an authorized user from an outside attack?
- What types of automation might aid in incident analysis?
- How do we know if we are under a sustained cyber-attack from a nation state? What might it look like?

2. INSIDER MISUSE MITIGATION

This workshop focus area concentrated on an assessment and user-validation of the research-oriented recommendations contained in the DoD Insider Threat Mitigation Report, that is set in the context of the overall insider threat to information systems. This threat is generally characterized as malicious actions by a disgruntled employee or an agent provocateur, disdain for security policy and practices, ignorance of security policy, and carelessness in applying security practices. The results of an August 1999 Santa Monica workshop (Anderson, 1999) outlined a variety of research and development directions of particular relevance to insider misuse are included in an Insider Threat Mitigation IPT Report. It was now felt to be time for users to review and discuss these recommendations and the plans for future research in this area. This session also addressed the following additional questions:

- To what extent are the IPT recommendations valid? What is missing?
- To what extent can mitigating insider misuse address the problem of detecting and reacting to external network attacks? Can insider misuse problems be used as a forcing function to help solve some other problems--such as the recent distributed denial of service attacks and information misuses by insiders?
- What other directions must be pursued to adequately address the anomalies, misuse, and malicious activity by insiders?
- What architecture and system design principles or characteristics would minimize the insider threat?

3. DEFENSIVELY ENGAGING THE ATTACKER

This focus area studied techniques that can be used to defensively engage an attacker. Issues included:

- How can we use active network defense mechanisms (deception, fishbowls, lures, etc.) provide information about the attacker? For example, is it useful to develop inferencing techniques to determine the sophistication of attacks, and therefore

understand the competence and possible intent of the attacker. If so, what can be done?

- Are cyber deception techniques useful in an attempt to confuse an attacker? Deception seems to be a central aspect of almost every military success, yet these techniques are not commonly used in information warfare scenarios.
- Is it useful to divert the attacker to an isolated system where he can be observed and more information can be gathered on the attacker's skills and intent? Is this type of attacker profiling useful in determining safe and appropriate response to attacks? Is it also useful in predicting the next moves of an attacker?
- What are the cautions and risks involved in using deception techniques and other means of active network defense?
- When should we use stealth, speed, what types of scenarios are worth a response, and what is a safe/appropriate response for a given scenario?
- What passive expert system traffic analysis can be used to help identify an attacker toolkit and skill level?
- A theme emerging from earlier workshops and reports is the need to provide "watermarking" of documents, data, etc. so that they can be traced. How can existing COTS software products be used or modified in a simple yet tamperproof way, so that the documents they produce (e.g., spreadsheets, slides, data files) can be tracked as they are transmitted, modified, stored, and retrieved within an organization? Can this be done not only in a way to aid analysis, but that will stand up in court as forensic data?
- How can such facilities be integrated into the actual operations of real world critical Defense information systems?

A three-day, invitational workshop was held, mixing plenary sessions with breakout sessions to explore the most promising aspects of the above topics in greater detail. The workshop also included a tour of relevant USCINCPAC facilities. Approximately forty invited

individuals participated, each involved with research in active network defense, or as future users of the proposed network defense capabilities.

Due to the sensitive nature of the third topic (defensively engaging the attacker), those discussions were held at the SECRET//NOFORN level; the other two breakout sessions were conducted at the unclassified level.

2. ANALYSIS OF INCIDENTS

This session discussed tools and techniques needed by operational users to aid in the analysis of computer/network attack incidents. The importance of this topic was underlined by site visits the first afternoon of the workshop to network operations centers within the Pacific Command, in which the status of networks throughout the command are continuously monitored and displayed, and operators attempt to discover anomalous probes and incidents within the mass of data being received.

OVERALL R&D PRIORITIES FOR ANALYSIS OF INCIDENTS

In intensive discussions between researchers and operators, the following R&D objectives were deemed to be of highest priority:

1. Reducing the "noise level" within the data, through automation
2. Better tools for correlating data from different sensors
3. Understanding how humans correlate intrusion detection system (IDS) information
4. Data abstraction: Creating higher-level abstractions to represent the "raw" data being received, for better human consumption and analysis
5. Means for updating the "signatures" of attacks so that they can be detected, in a manner similar to that used by virus detection tools
6. Tools for rapidly "profiling" a network: mapping its actual configuration
7. Determining what hacking tools were used in an attack
8. Maintaining the network state: "The first hack isn't enough." That is, techniques are needed to assure that network is resilient to any first attack--and to understand network status after that attack.
9. Metrics to determine the effects of potential response scenarios.

The discussion group focused on several of these themes in more detail, as indicated below.

REDUCING THE NOISE LEVEL WITHIN THE DATA

It was deemed vital that means be found of "reducing the noise level" in the data by less manpower-intensive means than are being used currently. In this context, "noise" results from the trivial hacker attacks and probes that are received by PACOM (and other DoD) networks daily. Noise also results from a large number of "false positives:" activity that trips an alarm, but is routine--for example, resulting from scans of the network by network administrators. Noise may also result from IDS system output that is determined to be irrelevant in light of present information. (However, that information may later be determined to be significant, so it should be available for later analysis and synthesis with other events, incidents, and data.)

How can the noise level be reduced? Themes discussed were:

- *Know your network*: Thorough knowledge of its topology and vulnerabilities can aid in distinguishing what is important from what is not.
- *Understand "normal" network traffic*. To isolate abnormal events, it is necessary to understand what constitutes normalcy within a network. Unfortunately, this will change depending on external events, such as exercises, world events, and so on.
- *Be aware of current hacker trends*. Most attacks use "scripts" provided on hacker bulletin boards, and discussed in "chat" sites. Familiarity with the operation of common hacker scripts can allow an analyst to determine its importance, how successful it is likely to be in uncovering or disabling network assets, and so on.
- *Update attack signatures in a timely manner*. As mentioned above, most attacks have a "signature" that will identify them, because they use standard scripts and protocols. If analysis tools can automatically check for such signatures within the raw data, they can abstract higher-level incidents out of that

data. This requires, however, that the database of signatures on which this process depends be updated often, because of the rapid evolution of attack scripts within the hacker community.

The group also discussed the importance of "patternless" intrusion detection--locating suspicious or anomalous patterns within the data that don't fit any predetermined signature or pattern. Such tools operate by building up and updating a real-time database of normal activity within a network, then alerting users to patterns of activity that fall outside some bounds of normalcy.

Another issue raised was the need to have "raw" data available for potential off-line analysis. Once an anomaly is detected, it is often important to look back in the data for precursor events, or larger patterns of which this incident is a part. This in turn raises the issue of the storage capacity needed, since gigabits of data pass through these network in seconds. Should absolutely raw data be stored, or is some higher-level summary sufficient? How long must it be stored? How can potential terabytes of stored data be accessed and searched for corresponding events?

The analysis time for some incidents is currently weeks, or even months. By greatly reducing the noise level in the data through the above means, and by providing powerful analysis and visualization tools, the aim is the reduce the analysis time to seconds or minutes.

AUTOMATING TASKS FOR NETWORK-BASED INTRUSION DETECTION SYSTEMS

The discussion focused on specific types of automation that are needed in intrusion detection systems. The principal types of automation deemed most important were these:

- *Filtering out noise* (without having to desensitize the sensor). See the discussion of noise in the previous subsection of this report. Based on existing attack signatures, definitions of normal behavior within the system, etc., the system should filter out this noise so that it doesn't overwhelm a manual analysis. (Since this filtering process will never be perfect,

however, the raw data should be retained for some period, in case analysis needs to delve into it for anomalous patterns masquerading as noise.)

- *Output consolidation for all IDSs used on a network.* A network of information systems may well contain a number of separate intrusion detection systems. It is important to consolidate their results, in order to see observe larger patterns. All IDSs should provide their output in a standard format (e.g., based on the XML format). There was discussion of the AIDE ACTD Translator, which reformats and consolidates IDS reports into a standard format. The consolidation of IDS outputs mentioned here is a subset of the larger problem of correlating data among different sensors. Tools are needed to aid in that process.
- *Data mining and visualization tools.* In general, a higher-level view is needed of the output of IDS systems than is currently available. Much richer visualization tools could provide the equivalent of a "flight simulator" or "cockpit" within which analyses might be conducted using the full pattern-interpretation capabilities of human analysts.
- *A template to generate reports based on IDS output.* It is necessary to report on anomalous behavior and incidents discovered within DoD networks. The burden of such reporting for an analyst would be reduced by providing templates and tools by which the generation of such reports could be automated to the extent possible.
- *Tighter, more timely, more accurate signatures.* The "looser" a stored signature of an event is, the more false positives and false negatives will result from its application. Such signatures must be as "tight" as possible to trigger alarms only when an event of the type sought is discovered in the data. And there must be highly timely means of updating these signatures, as mentioned earlier in this section.
- *Partitioning the database of IDS output.* It must be possible to perform flexible, tailored rearrangements, selections, and

partitions of the data resulting from IDS systems. For example, one might want to correlate certain types of suspicious activities with certain IP addresses. Such data selections and manipulations should be available to the analyst in a manner that allows a rich set of correlations to be explored.

- *A signature distribution system.* Once a new signature of an attack/probe method is created, it must be possible to distribute it quickly to all relevant IDSs within a command (or even within all DoD network operations centers).
- *Tools for traffic pattern analysis.* "Traffic analysis" refers to information that can be gleaned from an understanding of the patterns of data transmission, independent of the content of that data. For example, it would be noteworthy if numerous data packets were being sent to a specific foreign IP address from within a DoD network. Analysis tools are needed to aid in traffic analysis within DoD data networks, to help uncover anomalous patterns of activity worthy of further investigation.
- *The ability to add your own signatures to an IDS system.* Commercially available IDS systems, used by DoD, come with a set of signatures for anomalous behaviors to be detected. It is important that DoD sites have the ability to add their own signatures to such systems, both to keep them current with changing attack methods, and because some types of attacks (e.g., from computer-sophisticated agencies with foreign nations) may use probing and attack methods that differ from those normally used by hackers.
- *A tailorable, automated countermeasure.* A automated means of countering probes and attacks should be developed, but with the ability to manually configure it so that it may be tailored to the context and situation.

R&D FINDINGS AND RECOMMENDATIONS

The "Analysis of Incidents" discussion group concluded their deliberations with a set of "R&D Findings" which result from the

discussions outlined above. They are presented in five categories: those related to intrusion detection systems, network mapping/nodal analysis tools, anti-viral tools, damage assessment tools, and visualization tools.

Intrusion Detection Tools

The group felt that R&D should address and enhance current abilities to:

- *Evaluate IDSs.* Which have overlapping capabilities? Useful supplemental abilities? Which are easily tailored and configured? Which provide automated means for obtaining and installing the latest signatures? Etc.
- *Recommend which IDSs work best together.* What set of IDSs, used together, provides the best coverage? Can their outputs be easily correlated and synthesized into one overall report?
- *Disseminate information to the field.* As new capabilities are developed in IDSs, it is important that end-users "in the field" be aware of the state of the art, and that these developments are not just retained within a research community. What are the best mechanisms for providing this information dissemination on a continuing basis?
- *Develop an attack taxonomy.* Some attacks are variations on a theme of others. Some have fundamentally different approaches (e.g., a distribution denial of service, vs. a virus attack). A taxonomy of attacks should be developed, so that one can see quickly where a particular type of attack fits within the larger picture, and what the dangers and damages might be from a particular attack method.
- *Compare the attack space to IDS capabilities.* Given an attack taxonomy (see above) mapping the whole attack "space," how much of that space is covered by current IDS systems (and combinations thereof)? Where are the holes in IDS coverage?
- *Provide dynamic reallocation of IDS capabilities.* To support on-going missions, it is at times necessary to quickly

reallocate IDS capabilities. Remote installation of IDSs should be possible, and the ability to install mobile software "agents" to perform various detection and analysis activities.

- *Understand how humans correlate IDS information.* In what form should IDS output data be provided, to best capitalize on human correlation capabilities? In general, we need to understand human data correlation techniques better, in order to tailor our tools to these capabilities.
- *What non-IDS information can help identify attacks?* Attacks occur in a context--of world events; of hacker tools made available on various Internet bulletin boards; of U.S. exercises, deployments, and other activities. We should understand what information external to the network and its data events can help identify the types, sources, motivations, and perpetrators of attacks.
- *Provide relevant input to researchers.* In order for researchers in IDS systems to be maximally effective, they require inputs from the field to make their studies as realistic and useful as possible. Those inputs include:

  --Analysts' knowledge and expertise, as input to expert systems within the IDSs
  --Real-world data to support machine learning
  --Knowledge of current attack signatures
  --An operator's view of the IDS: What information needs to be displayed, in what amounts, what format, with what timeliness?

Network Mapping/Nodal Analysis Tools

Network protection, and incident analysis, is at times hampered by a lack of understanding of the topology of the network, and an understanding of the critical nodes for its operation.

The group felt it was important that network mapping and nodal analysis tools be developed and implemented, and that network operations personnel be trained to use them.

Examples of these tools and their usage known to the discussion group were:

- Communication Network Modeling Tool (CNMT) and Multi-Simulation Analysis Tool (being developed at the Applied Research Laboratories, University of Texas). These currently exist for telecommunications networks, and are being adapted for computer networks.
- Silent Runner, commercially available from Raytheon
- Bill Cheswick, of AT&T, performed a mapping of the Internet. The analysis includes a critical nodes study. He has a paper on this under development.

Anti-Viral Tools

The group recommended that an "intelligent" tool be developed for opening e-mail attachments to verify that they match standard file formats. Such a tool should operate in a user-friendly safe mode.

Damage Assessment Tools

A tool is needed that will evaluate damage from viruses and from actions of unauthorized users. It should determine what the virus is doing (e.g., copying files, transmitting information), and discover what actions an unauthorized user performed (e.g., data access, modifying files, etc.)

Visualization Tools

Much better visualization tools are needed to help analysts identify patterns of activity. A rich toolkit should be available to allow output from IDSs and other sensing mechanisms to be sorted, arranged, grouped and correlated, so that underlying patterns of interest might be detected.

3. INSIDER MISUSE MITIGATION

This breakout session focused on validating and prioritizing the key R&D recommendations contained in the DoD Insider Threat Mitigation Final Report of the Insider Threat Integrated Process Team (DoD-IPT 2000). To do this, they adopted the baseline definitions in the August 1999 Insider Workshop report (Anderson, 1999) for "insiders", "insider threats and vulnerabilities," and general characteristics of relevant protection, detection, and response technologies. The discussion focused on 24- to 36-month R&D recommendations, with the aim that all such recommendations were to be capable of being carried out, i.e., "actionable."

CRITERIA

In reviewing the recommendations in the DoD-IPT report, the following questions were posed as criteria:

1. To what extent are the IPT recommendations valid? What is missing?
2. To what extent can mitigating insider misuse address the problem of detecting and reacting to external network attacks? Can insider misuse problems be used as a forcing function to help solve some other problems, such as the recent distributed denial of service attacks and information misuses by insiders?
3. What other directions must be pursued to adequately address the anomalies, misuse, and malicious activity by insiders?

New recommendations resulting from group discussions were to address two issues:

1. What has, or would have, the greatest positive impact on your (the operational user's) environment?
2. What has, or would have, the greatest positive impact on the joint warfighting environment?

FINDINGS AND RECOMMENDATIONS

All IPT report recommendations reviewed were determined to be operationally relevant and significant, according to the prioritization criteria. The table below summarizes the highest priority, operational IPT report R&D recommendations, as determined by this breakout group's participants. Forty-six other IPT report R&D recommendations are considered important but of less priority when evaluated against the criteria. The numbers in the recommendation column of the table are those assigned to that recommendation in the IPT report. They are provided for ease of cross-reference. Recommendation numbers beginning with an "X" are new, created during this workshop and are not contained in the IPT report.

Priorities were determined in a two step process. First, the breakout participants identified nineteen recommendations that appeared to be most operationally relevant. Second, the participants further filtered the nineteen recommendations into three bands to determine which ones represented the most urgently needed solutions, requiring research and/or development, and where results could be achieved in the 24-36 month time frame. Banding results were:

Band One - 9 recommendations
Band Two - 6 recommendations
Band Three - 4 recommendations

Each Band One recommendation is described by the relevant operational issues or constraints, research objective, open research issues, success metric, and expected resulting product. The recommendations are listed in descending priority order.

Band Two and Three recommendations are not listed in priority order. The reader should not infer a priority based on the sequence of listed Band Two and Three recommendations in the table.

Some follow-up issues or questions resulting from this breakout session's discussions were:

- To R&D participants, the question, "What is the plan to get prototypes or sustainable product to the field?" must be addressed.
- Participants suggested that DoD set up an interoperability lab. CISCO offered to help set up such a lab.
- An industry participant suggested that DoD and/or NSA send a "top ten concerns" letter to CEOs of all security vendors.

The resulting ranking of recommendations arising from this session's deliberations are given in Tables 3.1 and 3.2, below.

Table 3.1

Band 1 Recommendations

Rank: 1	**Relative Rank:** 1	**IPT Recommendation: 4.5** Create technology providing a tamper-proof (detection and resistant) audit (data including network management stuff) trail recording the actions of individuals authorized access to sensitive data and networks
Operational Issues or Constraints: Tamper resistant vice tamper-proof audit trail is key research area here; differential access controls; tamper-proof probably not achievable; result of this R&D needs to work with deployed systems		
Research Objective: Develop fast integration technology that creates and preserves chain of custody using existing or new technologies. **Open research issues:** What is/constitutes a digital chain of custody? How do you achieve and maintain it sufficiently to support prosecution? What are the boundary conditions? How might these impact security classification issues (e.g., vulnerability data)? **Success metric:** Successful identification and prosecution of insiders (data will stand up in court); requires minimum additional sysadmin work load to implement the solution; audit trail is unalterable by insider **Product:** Media containing high assurance data integrity in the collection and maintenance of audit data (e.g., aircraft flight recorder-like device); needs to be solved for both NT and UNIX environments.		

Rank: 1	**Relative Rank:** 2	**IPT Recommendation: 6.3** Configure and deploy intrusion detection systems to monitor the activity of insiders
Operational Issues or Constraints: Feasibility of computer misuse/anomaly detection vice intrusion detection; sic, an insider is not an intruder		
Research Objective: Determine which IDS' (or combination thereof) are most effective in identifying and monitoring insider activity. **Open research issues**: How do you integrate data from multiple sensors? How do the sensors fail, and do they do it gracefully? How do you make sure the operator does not have or get information overload? Need standardized metrics to determine success factors/insight for this recommendation? **Success metric**: Identifying standardized metrics; multiple perspectives showing the same finding for the same case data; data must be right (providing a minimum of false positives and negatives); minimal impact of reporting back on available bandwidth. **Product**: A prototype interoperable, tool set/core capability which effectively monitors insider activity.		

Rank: 1	**Relative Rank:** 3	**IPT Recommendation: 6.4** Implement use of network mapping tools to detect alterations in the configuration of a network
Operational Issues or Constraints: R&D tools required to develop device mapping/configuration mapping tools vice implement network mapping tools; resources may constrain use/implementation; changes in scope if software components are in the class 'device'.		
Research Objective: Develop dynamic, near real-time configuration mapping tools for networks, systems and devices. **Open research issues**: Can autonomous agents be used to do this type of work? Need to determine the existing limits of what configuration alterations we can detect in devices and identify the detection response time. What is a significant alteration? What is an authorized alteration? What about detecting rogue devices? **Success metric**: Attribution to specific individual or device; response time; latency; false positives and negatives; minimal bandwidth impact **Product**: Real-time, Openview™-like capability that provides much greater granularity and an alarm to draw attention to alteration events.		

Rank: 1	Relative Rank: 4	**IPT Recommendation: 1.1** Develop and implement relevant definitions, methodologies and metrics tailored to the insider threat
Operational Issues or Constraints: Need to propose a standardized family of metrics, first; link to recommendation 1.3, within context of other IA metrics.		
Research Objective: Develop and standardize terminology, definitions, methodologies and metrics. **Open research issues**: What are useful, relevant metrics; meta-phenomenal nature of "insider" and interdisciplinary nature of research needed (social psychology, anthropology, computer sciences, mathematics, etc.). **Success metrics**: Commonly accepted meaning and value of terms, common methodologies - e.g., risk assessment (like the Service's Operational Risk Management), common tools (taxonomies, levels of severity and exposure (likelihood, reliability, visibility), etc.) and common metrics among CI, LE and cross-industry for a; real-time and retrospective validation; irrefutable indicators of security health/status (identification of value-added for operational decision-maker). **Product**: Prototype/framework taxonomy of terms, definitions and metrics.		

Rank: 1	Relative Rank: 5	**IPT Recommendation: 4.6** Consider means by which changes can be traced in all documents generated within an organization, by simple and tamper-proof (resistant) modifications to existing, widely used office automation programs
Operational Issues or Constraints: See 4.5 comment on tamper-proof issues.		
Research Objective: Develop a flexible tool for file "watermarking". **Open research issues**: How do you trace file history and movement, and changes to the file (genealogy? Associated non-repudiation? Identify key operational CONOPS constraints. **Success metric**: Demonstrable tractability of file types; minimal impact on BW. **Product**: User-transparent agent providing non-discretionary, tamper-resistant, application/OS independent, record of 'userid-made' changes to common file types.		

Rank: 1	**Relative Rank:** 6	**IPT Recommendation: 6.7** Develop better tools to detect, neutralize or eradicate the introduction of malicious "mobile code".
Operational Issues or Constraints: Significant R&D issue.		
Research Objective: Refer to developing DARPA malicious code mitigation program. **Open research issues**: Refer to developing DARPA malicious code mitigation program. **Success metrics**: prevention, detection, neutralization or eradication of malicious code. **Product**: Refer to developing DARPA malicious code mitigation program.		

Rank: 1	**Relative Rank:** 7	**IPT Recommendation: X.2** Develop a database of CND tools & their limits
Operational Issues or Constraints: Assessing security tools is of limited value unless centrally maintained site exists where the tool assessments can be shared. Moreover, a critical need exists to be able to select the 'right tool(s)' for the job because of the wide proliferation of security tools each with their own unique constraints, operational execution scenarios, and technological limitations dependent upon the technical task under consideration.		
Research Objective: Develop three capabilities in parallel: 1) a classification scheme that allows 'defenders' to establish the proper operating scenarios (and limitations) for each tool; 2) a Web-based capability to query, and update the classification results; and 3) a checklist approach that gives a decision tree so the defenders can pick the 'best/right' tool depending upon circumstances. **Open research issues**: Identifying incompatibilities and interoperability issues? **Success metric**: Testing done in an 'Underwriter's Laboratory' type model where skilled experts put each tool through a series of 'operational execution' steps that would allow the evaluator to determine: (a) Tool limits (technological, documentation, reliability, scale, utility, etc.). (b) Usage considerations (hardware limits, software interactions, software component limits, real-life usage results, when not to use the tool. (c) Complexity (how hard is it to use, does it require special training, typical mistakes the tool makes) (d) Interactions (how does this tool work with others; does the tool need other tools to work correctly) (e) Vulnerabilities (how does this tool fail? Gracefully? False positive/negative rates) **Product**: Database of tools and their limitations.		

Rank: 1	Relative Rank: 8	**IPT Recommendation: 1.5** Assess technologies currently available for dealing with the insider problem
Operational Issues or Constraints:		
Research Objective: To develop an integrated suite or taxonomy of modular existing, new and projected technologies for dealing with the insider threat ("Ask Jeeves for security" type decision tree tool). **Open research issues**: What existing, new and projected technologies are being or could be used for dealing with the insider problem? What characteristics of technologies make them applicable to the insider problem? **Success metric**: Understanding what existing, new and projected technologies can be applied to the insider. **Product**: A "Consumer Reports" database of existing, new and projected insider technologies.		

Rank: 1	Relative Rank: 9	**IPT Recommendation: 4.8** (Modified) Assess the vulnerabilities of wireless connectivity and remote system/network admin both require technological approaches, with relevant R&D issues.
Operational Issues or Constraints: Wireless (RF & IR) and remote admin access issues require R&D.		
Research Objective: Determine whether wireless technologies are connected to the network and identify resultant vulnerabilities if they are; include wireless requirements in security architectures. **Open research issues**: How do you detect the connection of wireless technologies within a network? What are non-forgeable identifier(s) of wireless devices? What is the role of physical and logical security measures such as physical interference (magnetometers, flooding the spectrum with noise), in countering wireless devices, enumeration of wireless technologies, and the vulnerabilities they introduce? What are the vulnerabilities of wireless devices? What are most cost-effective mitigation measures? **Success metric**: No known unknown wireless devices connected to a network. **Product**: Tool set for detecting wireless devices.		

Table 3.2

Bands 2 and 3 Recommendations

Rank: 2	**Relative Rank:** X	**IPT Recommendation: 1.3** Develop a database of insider events, characteristics and statistics.
Operational Issues or Constraints: Must link to and include CI and LE communities to ensure data is available.		

Rank: 2	**Relative Rank:** X	**IPT Recommendation: 1.10** Perform research on identifying critical information, automatically.
Operational Issues or Constraints: Aggregation related issue; on-going, beyond 24-36 month focus.		

Rank: 2	**Relative Rank:** X	**IPT Recommendation: 5.12** Continue research on developing a system security architecture sensitive to demands of the insider threat.
Operational Issues or Constraints: Fundamental R&D issue; Because its important to understand the taxonomy; this item is tied to recommendation X.3.		

Rank: 2	**Relative Rank:** X	**IPT Recommendation: 6.8** Create a comprehensive list of system and user behavior attributes that can be monitored to establish normal and abnormal patterns to enable anomaly and misuse detection.
Operational Issues or Constraints: Significant R&D issue.		

Rank: 2	**Relative Rank:** X	**IPT Recommendation: 7.3** Conduct research on means of reacting to suspected insider malicious activity.
Operational Issues or Constraints: Insider misuse, including malicious activity.		

Rank: 2	**Relative Rank:** X	**IPT Recommendation: X.3** Importance of architecture, authentication, defensive measures that transcend the insider, M&S of insider events, differential access control.
Operational Issues or Constraints:		

Rank: 3	Relative Rank: X	**IPT Recommendation: 4.7** Deploy a DOD Public Key Infrastructure (PKI).
Operational Issues or Constraints: How do you revoke keys and know it has been done? What are operational implications? Latency?		

Rank: 3	Relative Rank: X	**IPT Recommendation: 6.9** Establish a broad-based, long-term research program in anomaly and misuse detection addressing the insider threat.
Operational Issues or Constraints:		

Rank: 3	Relative Rank: X	**IPT Recommendation: 7.4** Conduct a long-range research program on reaction to insider threats.
Operational Issues or Constraints: Fundamental, base science R&D requirement and effort; will go beyond 24-36 month window.		

Rank: 3	Relative Rank: X	**IPT Recommendation: X.1** Need integrated plan that coordinates the near-, mid- and long-term R&D efforts.
Operational Issues or Constraints: Ties to recommendation 6.9.		

4. DEFENSIVELY ENGAGING THE ATTACKER

Information system security (INFOSEC) research has been quite successful in producing numerous, formidable safeguards of information, information systems, and information infrastructure. However, many significant deficiencies remain. As shown in Table 4.1, these can generally be described as relating to the reactive nature of INFOSEC--building walls and plugging holes against dimly-perceived external threats--and the chaotic, noisy, evolving character of the electronic environment.[1]

Table 4.1

Enduring Challenges to INFOSEC

	Challenge
1	INFOSEC generally cedes the initiative to the adversary.
2	INFOSEC usually innovates in response to an adversaries' *demonstrated* capabilities or *established* friendly shortcomings.
3	There are many potential attackers, of differing sorts.
4	The attackers may have many different motives, and many different objectives.
5	In virtually every case, attackers possess initial anonymity, and attacks emerge from under a blanket of secrecy.
6	There is a very large amount of data which *might* be relevant to the defense.
7	There is a very large amount of 'noise' surrounding the relevant data.
8	There are many 'locations' to defend, and adversaries may be insiders or outsiders.
9	Usually, the INFOSEC mission must be performed while keeping the protected systems up and running
10	Legal challenges may constrain (or even hamstring) optimal defense plans.

[1] Table 4.1 and some of the surrounding discussion is taken from "Employing Deception in Support of INFOSEC: An Ongoing Research Effort in RAND's National Defense Research Institute," by Scott Gerwehr. This paper was given as read-ahead material to this session's participants.

It is felt by some that INFOSEC research focuses perhaps too intently on traditional defensive measures (firewalls, encryption, biometrics, etc.) to the neglect of others. Another category of approaches could be characterized as defensively *engaging the attacker*. That approach was the focus of this breakout session's deliberations.

DECEPTION AS AN ACTIVE NETWORK DEFENSE MEASURE

One relatively neglected measure, which the historical record suggests is a potent tool on both offense and defense, is deception. Used effectively, deception is also a valuable mechanism of intelligence-gathering, which is a critical piece of the INFOSEC puzzle. An assumption discussed in this session was: *Deception in information systems, appropriately deployed and used, can address each of the INFOSEC challenges listed in Table 4.1.*

While often viewed as unsavory to practice and complicated to manage, there is little doubt that when employed successfully, deception is among the most powerful instruments of conflict. The discussion explored how deception might contribute to INFOSEC in two general ways:

- What protective value deception measures provide against a range of attacks on information infrastructure
- What protective value deception measures provide *which is not already provided by other types of defensive measures*.

As the latter category suggests, persons working in the field feel that deception offers *unique* defensive capabilities to INFOSEC. Moreover, deception is not a single tool: it is a diverse array of measures which may be employed individually or in depth, as simple schemes or complex ruses.

An active research program is underway at RAND on developing and testing a "deception toolkit" that would allow such measures to be introduced into computer networks, and for their effectiveness to be evaluated in a series of controlled experiments. One benefit of this workshop, bringing together researchers with field operators, was the opportunity to obtain feedback from operators on the likely importance

and value, and cautions and risks to be considered, in introducing such techniques into existing networks as active defensive and intelligence collection measures.

OTHER ACTIVE NETWORK DEFENSE MEASURES CONSIDERED

Other examples of active network defense measures discussed by this session were:

- Active network intrusion response systems, in which automated responses are triggered when an intrusion is detected
- Autonomic intrusion assessments, wherein assessments of the significance and danger of an intrusion are produced automatically by software agents and analysis programs
- Improving intrusion detection systems, and metrics for their evaluation.

As can be seen above, and throughout this session's deliberations, there was some overlap of discussion with that of the "Analysis of Incidents" session (see Section 2, above). Both sessions felt that their discussions must focus on intrusion detection and tools for automated analysis of incidents, because the outputs from those systems are vital both as aids to human analysis and judgment, and also as inputs to automated "active network defense" systems that react autonomously to detected intrusions and related events.

NEEDS TO BE ADDRESSED FOR ACTIVE NETWORK DEFENSE

The session listed a set of needs by field operators, and questions raised by techniques for "defensively engaging the attacker."

- How does one do battle damage assessment (BDA)? That is, if you perform certain active measures that "engage the attacker," how can you measure what effects those measures have had?
- What can be done about two dangers resulting from automation: (1) a false sense of security; and (2) removing the analyst from the underlying data? It is possible that a misplaced faith in the abilities of automated tools (that are, in fact,

missing key patterns in the data) may lull one into a sense of security. It is also possible that the analyst, given only synthesized, abstracted views of the data as a result of the operation of a set of automated tools, may lose a vital "feel" for what is happening at an IP packet level--a feel that operators said was vital in their current analyses.

- What are better ways of identifying the attacker? Unless we can be sure who is perpetrating the action (although that might be disguised by the use of various intermediate cutouts), there will be great (and appropriate) reluctance to "actively engage the attacker."
- How can we measure "normal" network behavior? (The importance and relevance of this was discussed in Section 2.)
- How can we locate single probes lost in the noise, which may be subtle mapping of our networks by a sophisticated adversary?

TOOLS AND CAPABILITIES AVAILABLE IN THE SHORT TERM FOR ENGAGING THE ATTACKER

In order to know what research is needed on defensively engaging the attacker, it is first incumbent to understand what capabilities will be becoming available in the relatively short term--e.g., within one year from the workshop (i.e., by July 2001). Participants listed the following activities that they knew were underway.

Deception Toolkit

RAND has a project underway that is developing specific tools and techniques for introducing deceptive measures into information systems. It is involved in a set of controlled experiments within a laboratory setting to test the effectiveness of those measures, and to develop metrics for such evaluations. It is developing a concept of operations (CONOPS) for their use, and considering counter-deception methods that might be used against these measures, or in detecting measures being used against our systems. Finally, it is considering deception as an intelligence gathering mechanism.

An initial framework for considering these issues is available in a limited-distribution document resulting from last year's effort (FY99);

see Gerwehr et al (1999). A similar document describing the results of this year's (FY00) work will be available in October 2000. Distribution is likely to be limited to U.S. government officials and government contractors working in this area.

Intrusion Detection Systems and Autonomic Response Mechanisms

Research programs are underway at MIT Lincoln Laboratory on intrusion detection systems and their evaluations, and on autonomic information assurance. Those programs were described by one of the participants.

IDS Evaluation. As part of their IDS evaluation program, they have synthesized a substantial amount (weeks' worth) of network traffic data, then annotated that data with labeled attacks. They have created quantitative, statistically relevant measurements for analyzing network packet data. This corpus of labeled data is in use by researchers and developers at more than 80 sites.

The data sets being developed this current year (2000) include inclusion of scenario-based attacks.

Mini-SimNet. Under sponsorship of the Air Force Electronic Systems Center, Lincoln Labs is developing a network simulation with a number of features:

- It has a graphical user interface (GUI), allowing much of the underlying complexity to be automated
- Traffic is generated "on the fly" in real-time
- It is configurable; the user can select the amount of traffic, the service to be attacked (e.g., ftp, telnet), and the type of attacks desired
- One can install scripts to configure hosts and services
- It has automated scoring and verification.

This environment will allow Lincoln to test new information assurance (IA) algorithms and systems, such as sensor fusion and correlation.

Probe detector. A Ph.D. thesis being completed at MIT involves the training of neural networks that store connection information about a particular network. It profiles normal network behavior to detect anomalous network traffic. The hope is that it will aid in detecting "slow probes" that might otherwise go undetected within normal network data streams.

Autonomic information assurance (AIA). As part of a larger DARPA information assurance program, Lincoln is studying dynamic, automated control of defenses and reflexive, automated responses to attacks in "machine-time" (e.g., sub-second). An AIA module might operate at many layers within a host computer system: NIC card, instrumentation, wrappers, component control, and the like. It performs sensing, detection, arbitration, and response activities, governed by coded policy statements, messages from other modules, and the raw data stream. It can send messages to other modules and commands to the computer system.

Data Correlation and Visualization

Session participants were briefed on activities underway within NSA's Active Network Defense Research group. These include:

- *Patternless intrusion detection*, based on highly novel ideas regarding the energy/entropy/temperature calculation for a network--providing a statistical overview of a network's "health"
- *Protocol validator (PV--"peavey").* This tool performs data reduction, provides playback of tcpdump files with filtering, detects protocol (RFC) violations, and allows the removal of "good" data to isolate more problematic data for further analysis.
- *Visualization* is provided by a tool called Propeller, that allows a user to visualize clusters of network activity.
- *Multisensor data correlation* is provided by a tool called Panda, that places multisensor data into a relational database

so that it can be selected and correlated using standard relational database techniques.

The above samples from RAND, MIT, and NSA are merely examples of currently- or soon-to-be-available tools, data sets, and techniques that might be employed for active network defense. The above listing is certainly not meant to be exhaustive. (For example, very significant amounts of work are being done within DARPA's Information Assurance program that aren't represented in the above overview.) However, these do give some indication of tools that might be employed--at least on an experimental basis--within operational network operations centers within the next year or so.

EFFECTS AND MEANS FOR DEFENSIVELY ENGAGING THE ATTACKER

To synthesize over two days' worth of discussions on defensively engaging the attacker, the session listed twelve specific effects that were desired, and for each listed various means by which those effects could be achieved. Some of this discussion necessarily paralleled that of the "Analysis of Incidents" group (see Section 2), because many of those analytic techniques and tools are required as precursors to an active engagement of the attacker.

In the listing below the effects desired are numbered, with the means to achieve them shown as subsidiary bulleted items.

1. Deny Enemy Access

This can be achieved by:

- *"Inoculating" a system*: Provide secure, automatic transmission of patches. Note that this requires an independent, secure channel for such transmission (because the network to be inoculated might be compromised). It is also dangerous if this capability is compromised or "gamed," so that it is turned into an automated means of introducing malevolent software within a system.

- *Limiting access to system and network information* (for example which operating system is being used; what network connections are in place). This can be achieved by hiding and other deceptive measures.
- *Dynamic reconfiguration of the network* to isolate an attacker. In this manner, an attacker might be "walled off" or otherwise isolated in real-time as his or her presence in a portion of the system is detected.

2. Limit Enemy Access

An unauthorized party's access may be limited by:

- *Confusion (by deceptive means).* If the underlying nature and assets of an information system can be disguised to the point where an attacker is confused as to its true architecture and assets, that alone may discourage the attacker from proceeding. Note that this implies that deceptive measures (e.g., ones hastily put in place) need not always present a consistent "story" to the attacker to have an effect.
- *Misinformation regarding the host.* If deliberate misinformation is provided to the attacker (e.g., regarding the type and version of the operating system in use), anyone acting on that misinformation signals a possibly hostile intent.

3. Create an Inability for an Enemy to Affect the System, Given Access

If an intruder gains some level of access to the system, we should consider active measures by which he or she is denied the ability to affect this system in deleterious ways. (No specific means were listed for achieving this desired effect.)

4. Learn the Attacker's Intent, Identification, and Level of Sophistication

This might be achieved through the following means:

- *Entice the attacker.* If the attacker remains for a considerable period of time within the system (e.g., by

investigating interesting "honeypot" files created and stored for that purpose), one might increase his or her latency within the system long enough for various tracing measures to be initiated, allowing one to learn the attacker's identity.

- *Use levels of protection or deception.* By employing levels of protection or deception, each of increasing subtlety or difficulty, one can gain a measure of the attacker's level of sophistication by discovering how many such levels he or she has broken through, or seen through.
- *Deflect or divert the attacker* to alternative sites with diverse types of information. This is a variation on the "honeypot" idea: By "pointing" the attacker at various diverse sites, one can gain an understanding of his or her intent and interests by learning which appear most appealing to him or her (e.g., judged by time spent at various sites, data files accessed, downloads initiated).

5. Detect the Presence of the Attacker

Means to detect the attacker should be quite familiar to a reader of this document. They include:

- *Automated tools to isolate anomalies.*
- *Visualization tools*, to aid a human in digesting large quantities of information, and finding subtle patterns within the data.
- *Use of "canaries" with the system*--attractive "bait" files that signal when they are accessed. These files should be designed, and normal system users trained, so that they are avoided in normal system operations. When such files are touched, they alert systems personnel that an intruder may be present.
- *Watermarking of data*, and checking for its presence at various bottlenecks, gateways, or other intelligence collection measures. If data generated within a facility (or organization, or CINC, ...) were digitally watermarked in an unobservable and ineradicable manner, the presence of that

watermark might be scanned for in abnormal places (e.g., on documents stored on hacker bulletin boards, in transmissions over various communication media). Its presence where it shouldn't belong would indicate a leakage of data from the facility or organization. The watermark may aid in tracking such leakages.

6. Mapping Internal Relationships and Organizations of Intruders

Through various traffic analysis, SIGINT, and other means it may be possible to map the communications relationships of intruders, and organizations within which they operate. No explicit means were given for accomplishing this desired effect.

7. Training of Our People

It became clear during the deliberations of this session that a continuing training program is vital for network operations personnel. Such a training program should teach methodologies and principles that will have lasting value, not specific tools that may quickly become obsolescent. Means for achieving this training include:

- *Use playback tools* (e.g., based on tcpdump files) that allow repeatable data patterns and sequences to be investigated and studied.
- *Create games that are played in a "cockpit"*--not unlike video games, to engage the interest of sysadmins and operators. That cockpit might in fact be similar to, or identical to, a richer analysis environment to be provided to operators for normal network data analysis.
- *Use teams of defenders vs. teams of attackers on a training network.* Such a "red" vs. "blue" team atmosphere creates interest, and may generate novel ideas regarding attacks and defenses.
- *Consider a network of honeypots as part of a training network.* Training of operators might include the creation and distribution of attractive "honeypot" files within a network, so help gauge the interests of an attacker.

8. Reduce the Amount of Data to Be Analyzed

As mentioned in Section 2 and earlier in this section, much remains to be done to provide automated aids to reduce the volume and type of network data to be analyzed in the search for anomalies.

9. Reduce the Noise in the Data

See Section 2 for a discussion of the importance of this measure, and means for achieving it.

10. Map and Understand Our Own Nets, Including Their "Backdoors"

Means for achieving this "understand thyself" dictum include:

- *Use commonly available mapping tools within DoD.* But it was remarked that relying on DoD-wide standards for data mapping causes delays and decided what those standards will be, and in their dissemination. Perhaps more local flexibility in choice of, and experimentation with, mapping tools is called for.
- *Check that our patches are installed.* Most attacks by outsiders exploit known, published vulnerabilities within operating systems and application packages. It is vital that patches to fix these flaws be made immediately after they become available. Part of understanding our own network situation would be automated scans to determine whether all relevant patches are installed and operational.

11. Find Low-Level Probes; Separate Them from the Noise

One means for achieving this goal was highlighted in the session's presentation from Lincoln Laboratory. That is:

- *Consider the use of neural net technology.* As mentioned earlier, there are promising indications that appropriate use of neural nets might "remove the time element" and allow the connection of events (e.g., low-level probes) separated in time.

12. Monitor the Actions and Reactions of an Attacker

The primary means to achieve this effect is:

- *Create an environment in which it is possible to rapidly configure and install a comprehensive, portable "fishbowl"* within the software system. Using these facilities, all the activities of an intruder can be monitored in real time. (This, of course, requires detection of the attacker in real time, using other techniques listed in this and other sections of this report.)

To summarize the deliberations of this session, it is by no means premature to concentrate on "engaging the attacker." However, doing so requires that a set of preliminary steps be undertaken first, such as putting various detection and analysis measures in place, and knowing your own critical nodes, links, and data files. It was widely believed that the various deception measures presented to this group were a useful addition to the armamentarium of an active network defender. Those measures must be used with caution, but they provide some capabilities for luring and engaging an attacker that are not otherwise available, and provide one of the few known means of assessing an attacker's intent and level of sophistication.

It was also realized that it is not necessary to identify the attacker in order to engage him or her in some useful ways. For example, it may be possible to "change his utility function" by making it appear sufficiently hard, or ambiguous, to obtain the desired information or effects. One might also divert the attacker (whoever he or she is) to less-vital assets.

The session concluded with a statement that these direct discussions between researchers and field operators were extremely valuable in helping to focus research activities, and in providing operators with information on tools and techniques that may be directly relevant and available in the not-to-distant future, for possible test and evaluation in operational environments.

5. CONCLUSIONS AND RECOMMENDATIONS

GENERAL ISSUES

Some topics emerged during the workshop as general issues or themes of general concern to the community:[2]

1. Prerequisites for Intrusion Detection

A recurring theme throughout the workshop was the need to know your network as a prerequisite to performing intrusion detection or in distinguishing between normal and abnormal traffic. The knowledge required included a mapping of the network (at both a high macro as well as a fairly detailed level), identifying the components of the network, their dependencies, usage, strengths, and vulnerabilities.

2. Determination of Damage and Attribution

Several questions arose regarding how to determine what information has been compromised when an attack has taken place. There was interest in whether analysis could determine not only the identity of the attacker, but the geolocation of the source of the attack, the hacking tools used and the intent of the attacker as well.

3. Data Management

Data management was described as being very user-dependent. For the "master analyst," sophisticated analysis tools were of interest. However, the experienced analysts at the meeting explained the need to always be able to review the raw data.

For the less experienced analyst, data support via data reduction tools were described as essential. Intelligent data reduction was described as noise reduction in order to find the "nuggets" of useful information in the large volume of data produced by intrusion detection systems.

---

[2] This section is adapted from notes taken by participant Sara Matzner.

Data abstraction is needed to correlate events across time and space. Obstacles to data abstraction that were mentioned included the numerous different data output formats used by different sensors and IDS systems.

4. Tools for Intrusion Detection

In discussing the tools used for data analysis, it was noted that where the tool was located on the network and how the analyst used that particular tool was as important as the specific tool that was used.

Tools were employed for several purposes:

To perform statistics gathering for reports
To perform real time intrusion detection
To perform post-analysis

Several discussions revolved around the issue of whether the tools used by a hacker are any different that those used by a more sophisticated user with a more insidious intent.

A list of the specific tools for intrusion detection in use by persons participating in the workshop includes:

JIDS along with SNORT and SURVEY
ASIM
NetRanger.

When discussing GOTS vs. COTS IDS, the need to be able to customize COTS products for specific government requirements was described as a major concern.

5. Survivability

A note of concern was expressed about the fact that we are using the same networks for notification as for normal traffic. Also noted was the fact that some sensors maintain network state better than others.

6. Expectation Management

Issues raised were: How should we manage expectations and needs of government in this area? How can researchers help with short term needs and still develop and plan for the long term?

Directions for Future Research

The following list describes the group's overall assessment of directions for needed research in the area of intrusion detection.

- Rapid network profiling for a dynamic environment, including tools to provide a macro-view of a network.
- Data reduction and data correlation tools.
- Real world data sets with ground truth. This data needs to be coupled if possible with access to the analyst involved in the reporting of the incident.
- Honeypots that are deliberate bait files that would signal when they are touched.
- Establishing a consistent format for the reporting from intrusion detection systems.
- Automatic updates for patches.
- Visualization tools for real time display, and an executive summary graphic.
- A taxonomy of host-based and network-based attacks that is widely distributed to the community.
- Dynamic reallocation defensive assets of IDS database capabilities.
- Customizing GOTS IDS to be tiered on top of COTS IDS.

RECOMMENDATIONS

Several overall recommendations result from the workshop.

- It was agreed that researchers should work toward establishing a consistent format for reporting from intrusion detection systems. Cisco systems is interested in XML development of output from sensors. At the present, there is no mandatory output format for an IDS.

- Industry representatives stressed the need to find a means (at a significantly high level of authority) to let industry know the needs of the government.
- Several participants agreed to establish an informal exchange of information among analysts.

Overall, it was felt by all participants that the workshop was very useful. The agenda was ambitious, but most of the objectives were met. The gap between operational users of intrusion detection technologies and the advanced technologies being pursued by the research community is significant. One end of this gap (at the operational user end) primarily involves implementation and resources issues: tools and techniques exist today that could significantly help the analysts, but they are not being identified and provided, and the analysts are not receiving the necessary training on these tools. This situation provides many opportunities for initiatives that might be undertaken by sponsors of this research.

It is clear that analysts at PACOM, especially within DISA-PAC, are willing to share their data and domain knowledge with the research community. Steps should be initiated to take advantage of this offer.

It became very clear that PACOM analysts are unable to completely analyze the voluminous output of their intrusion detection systems, and are desperately in need of automation. In some cases, this would require no more than a simple PERL script, which they could theoretically write for themselves if they had the time; however, the operational requirements of their jobs precludes it. This situation presents a significant opportunity for the research community, that might follow the following model:[3]

---

[3]Readers may well note similarities between this description and the "spiral model" of software development pioneered by Barry Boehm. See, for example, its classic formulation in Boehm (1988).

1. Analyst (*A*) presents researcher (*R*) with a (hopefully) well-defined problem, the solution to which would be immediately useful. This might entail *R* spending some time interacting with *A*, familiarizing himself with the data.
2. *R* writes a (perhaps fairly simple) script to address the problem.
3. *A* uses the script, and comes back to *R* with proposed improvements. This is almost always the case in software development--the end users don't really know what they want until they begin to see the possibilities regarding what can be done.
4. Iterate on 2 and 3.

Both *R* and *A* benefit from this process: *A* gets something that makes his life easier, and *R* gets to observe the phenomenon in depth and become familiar with the real problems from the insider's perspective. *R* then uses his knowledge of related fields (e.g., operations research, probability and statistics, information retrieval) to come up with solutions to deeper and more difficult problems.

A number of specific follow-up activities resulted from contacts made during the workshop, which will be pursued by various of the participants.

Appendix

A. WORKSHOP AGENDA

DAY 1 - JULY 11

0730	Registration/Breakfast *(Imperial Suite)*		
0800	Plenary: Welcome; Workshop overview Dick Brackney *(Imperial Suite)*		
0820	Plenary: Sponsor's Opening remarks Tom Bozek, OASD/C3I LtCol Joe Squatrito USAF, USCINCSPACE J39 Jerry Hamilton, DARPA Dick Brackney, NSA/IARO *(Imperial Suite)*		
0900	Video: DARPA IA&S Strategic Cyber Defense Vision - Jerry Hamilton Video: Infrastructure Surety Program II, NRO, 2/15/00 (U) - Tom Bozek *(Imperial Suite)*		
0940	Plenary: CINCPAC J39 DIO briefing Harvey Blacker, NCPAC		
1010	BREAK *(Imperial Suite)*		
1030	PAC-CERT briefing John Coller, DISA-PAC *(Imperial Suite)*		
1100	Breakout 1a **Analysis of Incidents** Chair: Dick Brackney *(Imperial Suite)*	Breakout 2a **Insider Misuse Mitigation** Chair: Tom Bozek *(Kauai Room)*	Breakout 3a **Defensively Engaging the Attacker** Chair: Bob Anderson *(Molokai Room)*
1230	LUNCH *(Waikiki Suite)*		
1330	Guided tour of USCINCPAC CERT, NOC, other local facilities Harvey Blacker, NCPAC		
1830-2000	Reception *(Australia Suite)*		

DAY 2 - JULY 12

0730	Breakfast *(Imperial Suite)*		
0800	Plenary: Breakout session status report *(Imperial Suite)* What are key issues? What approach should be pursued during remainder of workshop?		
0845	Breakout 1b *(Imperial Suite)*	Breakout 2b *(Kauai Room)*	Breakout 3b *(Molokai Room)*
1000	BREAK *(in breakout rooms)*		
1020	Breakout 1c *(Imperial Suite)*	Breakout 2c *(Kauai Room)*	Breakout 3c *(Molokai Room)*
1200	LUNCH *(Waikiki Suite)*		
1330	Breakout 1d *(Imperial Suite)*	Breakout 2d *(Kauai Room)*	Breakout 3d *(Molokai Room)*
1530	BREAK *(in breakout rooms)*		
1545	Plenary: Interim results from breakout sessions *(Imperial Suite)*		
1800	Dinner Speaker: LTC Tim Gibson, Chief PACOM Network Security Division *(Australia Suite)*		

DAY 3 - JULY 13

0730	Breakfast *(Imperial Suite)*		
0800	Plenary Speaker: Col. Rakestraw, Deputy J6, PACOM *(Imperial Suite)*		
0845	Plenary: Final guidance to breakout sessions *(Imperial Suite)*		
0900	Breakout 1e *(Imperial Suite)*	Breakout 2e *(Kauai Room)*	Breakout 3e *(Molokai Room)*
1030	Break *(in breakout rooms)*		
1045	Breakout 1f *(Imperial Suite)*	Breakout 2f *(Kauai Room)*	Breakout 3f *(Molokai Room)*
1200	Lunch *(Waikiki Suite)*		
1330	Plenary: Reporting by breakout groups of findings, results, recommendations *(Imperial Suite)*		
1445	Break *(Imperial Suite)*		
1500- 1630	Plenary: Further reporting by breakout groups; Summary and conclusions by workshop sponsors *(Imperial Suite)*		

B. WORKSHOP PARTICIPANTS

Sponsors

Joseph Squatrito, LTC	SPACECOM, J39	(n/a)
Mike Skroch	DARPA	mskroch@darpa.mil
Dick Brackney	NSA/IARO	RCBrackney@aol.com
Tom Bozek	OASD/C3I	tom.bozek@osd.mil
Larry Merritt	DoD, Ft. Meade MD	l.merrit@radium.ncsc.mil

Support

Harvey Blacker	NCPAC	hdblacker@ hawaii.rr.com
Michael Rohrer	BBN	mrohrer@bbn.com
Bob Anderson	RAND	Robert_Anderson@rand.org
Peter Neumann	SRI International	Neumann@CSL.sri.com
Virginia Kerry	SPACECOM NSA Rep, J39	gtkerry@aol.com

Breakout Session 1: Analysis of Security Incidents

Dick Brackney, Chair	NSA/IARO	RCBrackney@aol.com
John Burke	DoD	JohnCBurke@aol.com
John Coller	DISA-PAC	coller@pixi.com
Michael Crabtree	DISA-PAC, PAC-CERT	crabtreem001@hawaii.rr.com
Larry Frank, Col.	JTF-CND J3	frankl@jtfcnd.ia.mil
Tim Gibson, LtCol	USCINCPAC J67	
Tom Hetherington	U. Texas ARL	tomh@arlut.utexas.edu
Keith Konen	JICPAC	
Gigi LaTorre-Couch	PACAF RIOC	gigi.latorre@cidss.af.mil
Scott Lewandowski	MIT Lincoln Lab	scl@sst.ll.mit.edu
Sara Matzner	U. Texas ARL	matzner@arlut.utexas.edu
Tom McLaughlin	Hq USCINCPAC J6 TCCC	j6tcccio@hq.pacom.mil
Christopher Mellen	NCIS	cmellen@ncis.navy.mil
Jack Miller	Pacific RCERT	millerje@shafter.army.mil
Susan Ogawa	FBI	suhodges@leo.gov
Kevin Pardue	RIOC, PACAF	kevin.pardue@cidss.af.mil
William Rybczynski, GySgt.	US Marine Forces Pac	RybczynskiWH@mfp.usmc.mil
Jamie Turner	NCIS	JTurner@ncis.navy.mil
Stephen Vanden Bosch	BBN Technologies	swvanden@bbn.com

Breakout Session 2: Insider Misuse Mitigation

Tom Bozek, Chair	OASD/C3I	tom.bozek@osd.mil
Lee Curto, Capt.	JICPAC	curtolil@jicpac.pacom.mil
Tom Goldring	DoD	tgo@tycho.ncsc.mil
Jerry Hamilton	DARPA/SETA	ghamilton@schafercorp-ballston.com
Wayne Meitzler	Pac. Northwest Natl. Lab.	wayne.meitzler@pnl.gov
Peter Neumann	SRI International	Neumann@CSL.sri.com
Thomas Ray	US Marine Forces Pacific	rayta@mfp.usmc.mil
Janice Reitzell	NCIS	jreitzel@ncis.navy.mil
Deborah Rocco	NCIS	DRocco@ncis.navy.mil
Ron Schmucker	LLNL IOWA Center	rfs@llnl.gov
Scott Schulle	NSSOC	schulles@shafter.army.mil
Steven Sonnenberg	DARPA/SETA	steve.sonnenberg@avenuetech.com
Eric VonColln	SPAWAR	voncolln@spawar.navy.mil
Brad Wood	SRI International	bjwood@sdl.sri.com
Kevin Ziese	Cisco Systems	ziese@cisco.com

Breakout Session 3: Defensively Engaging the Attacker

Bob Anderson, Chair	RAND	Robert_Anderson@rand.org
Kris Axberg	HQ PACAF/DOIO	kris.axberg@cidss.af.mil
Harvey Blacker	NSA/NCPAC	hdblacker@hawaii.rr.com
Steve Corrigan, Capt.	JICPAC	ABcorrsj@pacom.osis.gov
Scott Gerwehr	RAND	Scott_Gerwehr@rand.org
Kevin Hutchison	US Marine Forces Pacific	
Virginia Kerry	USSPACECOM	gtkerry@aol.com
Larry Merritt	NSA/X	l.merrit@radium.ncsc.mil
John Moffett	NSSOC	moffettj@shafter.army.mil
John Pericas, LtCol.	USSPACECOM	(n/a)
Michael Rohrer	BBN	
Lee Rossey	MIT Lincoln Lab	lee@sst.ll.mit.edu
Ray Roy	DoD	reroy@alpha.ncsc.mil
Bryan Schlather	FBI	bschlather@fbi.gov
Joseph Squatrito, LTC	USSPACECOM	(n/a)
Karina Tam	NSA/NCPAC	kntam000@hq.pacom.mil
Scot Taylor, MSgt.	Hq, PACOM	sktaylor@hq.pacom.mil
Terry Woodhouse	MITRE	tjw@mitre.org

REFERENCES

1. Anderson, R.H., 1999. "Research and Development Initiatives Focused on Preventing, Detecting, and Responding to Insider Misuse of Critical Defense Information Systems: Results of a Three-Day Workshop." RAND CF-151-OSD, 1999.

2. Boehm, B.W., 1988. "A Spiral Model of Software Development and Enhancement," *IEEE Computer*, May 1988, pp. 61-71.

3. Gerwehr, S., J. Rothenberg, R.H. Anderson (1999). "An Arsenal of Deceptions for INFOSEC." RAND, Unpublished document, October 1999.

4. Insider Threat Integrated Process Team, Department of Defense (DoD-IPT), 2000. "DoD Insider Threat Mitigation" U.S. Department of Defense, 2000. Available at http://www.c3i.osd.mil